



Additional AMD Training Materials

**Presenter: Bob Robey
AMD @ HLRS
May 8th, 2025**

AMD 
together we advance_

Lots of different resources on AMD Hardware and Software

-
1. [Training Offerings](#)
 2. [Lab-notes \(short technical articles\)](#)
 3. [Videos from EPCC](#)
 4. [Videos and Slides from 2024 HLRS](#)
 5. [Videos and Slides from OLCF](#)
 6. [Books](#)

Training Offerings -- HPC and AI Training Events

We typically offer the following training types.

- **System Administration Introduction** – 1-2 hours for key staff at a site
 - Model installation repo, useful for application developers as well: <https://github.com/amd/HPCTrainingDock>
- **Training Workshop** – generally around 4 half-days virtual workshop with hands-on exercises. Up to 90 participants.
- **Hackathon** – typically a 3 full-day on-site event for around 8 teams. Often there is a pre-hackathon virtual workshop the week before the event to get teams ready for hackathon
- **Hackathons** are generally scheduled **six months after a Training Workshop** to allow applications some time to get on the system and work through their initial porting efforts.
- **Training events are customized** to the typical workloads and applications that are run at each site.

Contact for more training and outreach support

- **Bob Robey**, AMD Data Center GPU Global Training Lead Bob.Robey@amd.com
- **Giacomo Capodaglio**, AMD Data Center GPU Global Training Team Giacomo.Capodaglio@amd.com

AMD lab notes and ROCm blogs

<https://gpuopen.com/learn/amd-lab-notes>,
<https://rocm.blogs.amd.com/index.html>



Technical blog post series covering:

- Lessons learned from tuning a wide range of applications, libraries, and frameworks
- Implementations of computational science algorithms such as PDE discretizations, linear algebra, solvers, and more
- Instructions, guidance, and references on using libraries and tools from the ROCm software stack
- Best practices for porting and optimizing both HPC and AI applications
- Monthly release cadence

Home » Blogs » AMD lab notes

Originally posted November 14, 2022

AMD lab notes

Computational and Data science have emerged as powerful modes of scientific inquiry and engineering design. Often referred to as the “third” and “fourth” pillars of the scientific method, they are interdisciplinary fields where computer models and simulations of physical, biological, or data-driven processes are used to probe, predict, and analyze complex systems of interest. All of this necessitates the use of more computational power and resources to keep up with increasing scientific and industrial demands. In order to fully utilize emerging hardware designed to tackle these challenges, the development of robust software for high-performance computing (HPC) and Machine Learning (ML) applications is now more crucial than ever. This challenge is made even greater as hardware trends continue to achieve massive parallelism through GPU acceleration, which requires the adoption of sophisticated heterogenous programming environments and carefully tuned application code.

In this “AMD lab notes” blog series, we share the lessons learned from tuning a wide range of scientific applications, libraries, and frameworks for AMD GPUs. Our goal with these lab notes is to provide readers with the following:

- AMD GPU implementations of computational science algorithms such as PDE discretizations, linear algebra, solvers, and

GitHub repository: <https://github.com/AMD/amd-lab-notes>

Courtesy: Justin Chang, AMD

List of Lab-notes and Blogs

[AMD matrix cores](#)

[Finite difference method - Laplacian part 1](#)

[Finite difference method - Laplacian part 2](#)

[Finite difference method - Laplacian part 3](#)

[Finite difference method - Laplacian part 4](#)

[Introduction to profiling tools for AMD hardware](#)

[AMD ROCm™ installation](#)

[AMD Instinct™ MI200 GPU memory space overview](#)

[Register pressure in AMD CDNA™2 GPUs](#)

[GPU-aware MPI with ROCm](#)

[Jacobi Solver with HIP and OpenMP offloading](#)

[Sparse matrix vector multiplication - part 1](#)

[Affinity part 1 - Affinity, placement, and order](#)

[Affinity part 2 - System topology and controlling affinity](#)

[C++17 parallel algorithms and HIPSTDPAR](#)

[Application portability with HIP](#)

[Reading AMD GPU ISA](#)

[MI300A - Exploring the APU advantage](#)

[Seismic stencil codes - part 1](#)

[Seismic stencil codes - part 2](#)

[Seismic stencil codes - part 3](#)

[Creating a PyTorch/TensorFlow code environment on AMD GPUs](#)

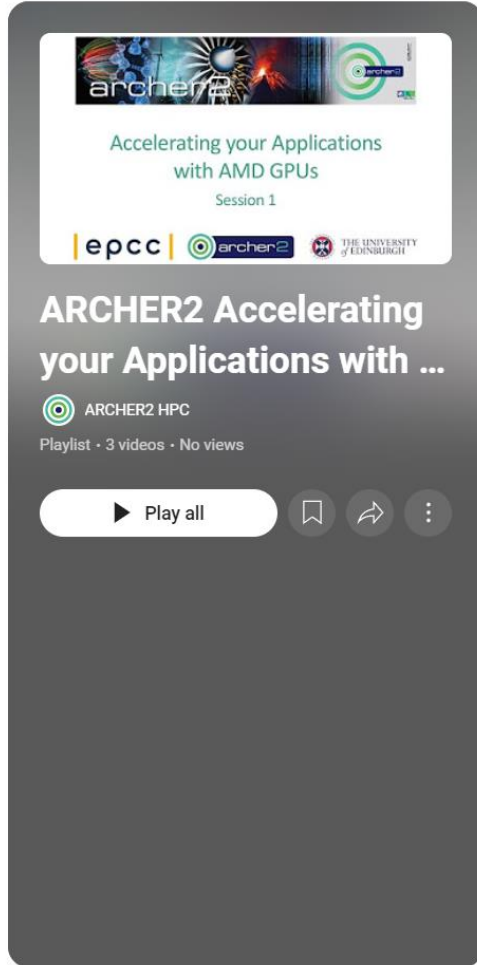
[Introducing AMD's Next-Gen Fortran Compiler](#)

[Graph analytics on AMD GPUs using Gunrock](#)

[Introducing ROCprofiler SDK - The Latest Toolkit for Performance Profiling](#)

Videos from EPCC

- Video recordings for this event are already available on YouTube:
<https://www.youtube.com/playlist?list=PLB4tvLCynFjQq2rOljEy39IDaugXcFPNa>



- 1 ARCHER2 Training: Accelerating your Applications with AMD GPUs (day 1)
ARCHER2 HPC · 31 views · 2 days ago
Accelerating your Applications with AMD GPUs
Session 1
2:48:33
- 2 ARCHER2 Training: Accelerating your Applications with AMD GPUs (day 2)
ARCHER2 HPC · 13 views · 1 day ago
Accelerating your Applications with AMD GPUs
Session 2
2:13:30
- 3 ARCHER2 Training: Accelerating your Applications with AMD GPUs (day 3)
ARCHER2 HPC · 7 views · 3 hours ago
Accelerating your Applications with AMD GPUs
Session 3
2:33:59

Videos and Slides from 2024 HLRS

- High-Performance Computing Center Stuttgart (HLRS) Stuttgart <https://fs.hlrs.de/projects/par/events/2024/GPU-AMD/>
April 2024

AMD Instinct™ GPU Training

Materials for this course



Day 1

- 01. AMD Presentation Roadmap.pdf
— [Slides](#), [Recording*](#)
- 02. Overview of Programming Model for MI200 and MI300 Series
— [Slides](#), [Recording](#), [QA & Exercises](#)
- 03. Introduction to OpenMP® Offload on AMD GPUs
— [Slides](#), [Recording](#), [QA & Exercises](#)
- 04. Real-World OpenMP® Language Constructs
— [Slides](#), [Recording](#), [QA & Exercises](#)
- 05. Advanced OpenMP®
— [Slides](#), [Recording](#), [QA](#)

Day 2

- 06. Introduction to Programming GPUs with HIP
— [Slides](#), [Recording](#), [QA & Exercises](#)
- 07. Porting Applications to HIP
— [Slides](#), [Recording](#), [QA & Exercises](#)
- 08. Optimizing HIP Applications
— [Slides](#), [Recording](#), [QA & Exercises](#)
- 09. Hip and OpenMP® Interoperability
— [Slides](#), [Recording](#), [QA & Exercises](#)

Day 3

- 10. Kokkos and Performance Portable Languages
— [Slides](#), [Recording](#), [QA & Exercises](#), [QA](#)
- 11. C++ Standard Parallelism on AMD GPUs
— [Slides](#), [Recording](#), [QA & Exercises](#)
- 12. GPU-Aware MPI with ROCm™
— [Slides](#), [Recording](#), [QA & Exercises](#)
- 13. ML and AI on AMD GPUs
— [Slides](#), [Recording](#), [QA & Exercises](#)

Day 4

- 14. AMD Debugger: ROCgdb
— [Slides](#), [Recording](#), [QA & Exercises](#)
- 15. GPU Profiling - Performance Timelines
— [Slides](#), [Recording](#), [QA & Exercises](#)
- 16. Kernel Profiling with Omniperf
— [Slides](#), [Recording](#), [QA & Exercises](#)
- 17. Additional AMD Training Materials
— [Slides](#), [Recording](#)

Documentation

- [ROCm Documentation](#)
- [ROCm Blogs \(search for HPC tag\)](#)
- [GPUOpen AMD lab notes](#)

Exercises and updates

- [Exercise document \(updates for Omniperf\) pdf](#)
- [Launch Parameter Tuning \(pdf\)](#)
- [LDS Occupancy Limiter \(pdf\)](#)
- [Register Occupancy Limiter \(pdf\)](#)
- [Strided Data Access Patterns \(pdf\)](#)
- [Algorithmic Optimizations \(pdf\)](#)

Additional slides

- [Affinity](#)
- [AMD Matrix Cores](#)
- [OpenMP with Fortran, and OpenACC](#)

Videos – Oak Ridge Leadership Computing Facility (OLCF) Training Archives

Search “AMD” in this webpage to find the **latest** talks: https://docs.olcf.ornl.gov/training/training_archive.html

2023-02-17	GPU Profiling	Alessandro Fanfarillo	Frontier Training Workshop	(slides recording)
2023-02-15	EPYC CPU and Instinct GPU	Nick Malaya, AMD	Frontier Training Workshop	(slides recording)
2022-10-11	Hierarchical Roofline Profiling on AMD GPUs	Noah Wolfe (AMD) and Xiaomin Lu (AMD)	Special Session	(slides recording)
2022-08-23	Understanding GPU Register Pressure	Alessandro Fanfarillo (AMD)	Special Session	(slides recording)
2021-05-26	ROCgdb and HIP Math Libraries	Justin Chang (AMD)	HIP Training Workshop	(slides exercises recording)
2021-05-25	Converting CUDA Codes to HIP	Julio Maia (AMD)	HIP Training Workshop	(slides exercises recording)
2021-05-24	Introduction to GPU Programming	Gina Sitaraman (AMD)	HIP Training Workshop	(slides exercises recording)
2021-05-20	MI100 GPU	Nick Malaya (AMD)	Spock Training	(slides recording)
2021-05-20	State of HIP	Nick Malaya (AMD)	Spock Training	(slides recording)
2021-05-20	Node-Level Profiling	Julio Maia (AMD)	Spock Training	(slides recording)
2019-09-06	Intro to AMD GPU Programming with HIP	Damon McDougall, Chip Freitag, Joe Greathouse, Nicholas Malaya, Noah Wolfe, Noel Chalmers, Scott Moe, Rene van Oostrum, Nick Curtis (AMD)	Intro to AMD GPU Programming with HIP	(slides recording)

Videos – continued

- ENCCS recorded our training sessions from our workshop for them at the end of November 2023. The full set of training materials is at the ENCCS (Euro National Competence Centre Sweden) including exercises, slides and videos -- <https://enccs.github.io/amd-rocm-development>
 - [Introduction to HIP Programming](#)
 - [Porting Applications to HIP](#)
 - [Getting Started with OpenMP® Offload Applications on AMD Accelerators](#)
 - [Developing Fortran Applications: HIPFort, OpenMP®, and OpenACC](#)
 - [Exercises](#)
 - [Architecture](#)
 - [GPU-Aware MPI with ROCm™](#)
 - [AMD Node Memory Model](#)
 - [Hierarchical Roofline on AMD Instinct™ MI200 GPUs](#)
 - [Affinity – Placement, Ordering and Binding](#)
 - [Profiling and debugging](#)
 - [OpenMP® Offload Programming](#)
 - [Introduction to ML Frameworks](#)
 - [Summary and outlook](#)

Videos – continued

- Pawsey Supercomputing Research Centre, Perth Australia
- [Pawsey AMD Profiling Video List](#)
 - [Introduction to ROCm Profiler - Day 1 Part 1](#)
 - [Introduction to OmniTrace - Day 1 Part 2](#)
 - [Hands-on workshop- AMD Profiling workshop - Day1 Part 3](#)
 - [Introduction to Omnipperf - Day 2 Part 1](#)
 - [Introduction to Roofline - Day 2 Part 2](#)

AMD software documentation

- ROCm documentation site is often the first resource (<https://rocm.docs.amd.com>)
- AMD Developer Site – <https://amd.com/developer>
 - Developer Site Videos: <https://www.amd.com/en/developer/resources/rocm-hub/training-videos.html>
- There is also now a book on HIP programming <https://www.barnesandnoble.com/w/accelerated-computing-with-hip-yifan-sun/1142866934>

- Oak Ridge National Laboratory has some excellent materials they have created from our briefings and their excellent staff:
 - [Crusher quick-start guide](#)
 - [Frontier user guide](#)

Books

Accelerated Computing with HIP – out of AMD Education Outreach group.



<https://community.amd.com/t5/rocm/accelerated-computing-with-hip-textbook/ba-p/620840>

New OpenMP GPU Programming book – **Programming Your GPU with OpenMP**, Tom Deakin and Tim Mattson

(2023) <https://www.amazon.com/Programming-Your-GPU-OpenMP-Performance/dp/0262547538/>

The best broad cover of parallel and high performance computing topics is the book **Parallel and High Performance Computing**, Manning Publications, by Robert Robey and Yuliana Zamora.

At <https://www.manning.com/books/parallel-and-high-performance-computing>



Disclaimer

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED 'AS IS.' AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

AMD, the AMD Arrow logo, ROCm, Radeon, CDNA, Instinct, EPYC and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

The OpenMP name and the OpenMP logo are registered trademarks of the OpenMP Architecture Review Board

© 2024 Advanced Micro Devices, Inc. All rights reserved.

AMD 